The Fast Multipole Method

The Fast Multipole Method provides a fast and precise approximation of n-body interactions. The algorithm scales linearly, and depends logarithmically on the inverse of the tolerance.

1. Introduction

Throughout this lecture, it might be easier to consider the Fast Multipole Method as it relates to evaluating electric potentials (a common problem in physics). Given a set of electric charges, each charge exerts a force on every other charge, and this force can be represented as an electric potential. We would like to determine the electric potential that is acting on every charge. Let $\{x_i\}_{i=1}^n$ be a set of n points in $\mathbb{R}^d$ (for our example, we can consider only the case where $d = 2$). Let A be an $n \times n$ matrix with entries $A(i, j) = G(x_i, x_j)$, where G is a kernel function. Our objective is to evaluate $u = Aq$.

Here, u is a vector containing the electric potential acting on each charge, and q is the vector of charges. $G(i, j)$ describes the potential acting on charge j due to charge i, and it can be written as

$$G(x, y) = \begin{cases} \log(x - y) & x \neq y, \\ 0 & x = y. \end{cases}$$

Naïvely, calculating u would require the evaluation of n^2 interactions. However, the Fast Multipole Method allows us to approximate the solution within an error of ϵ using only $\mathcal{O}(Cn)$ operations, where $C = \mathcal{O}(\log(\frac{1}{\epsilon})^2)$ (in two dimensions).

2. Establishing the FMM

In general, the kernel G is difficult to compute exactly. For our example from electrostatics, G is the solution to a partial differential equation. The Multipole Expansion gives an approximate "separation-of-variables" solution to the PDE, and we can write this solution as a sum of matrix products:

$$G(x, y) \approx \sum_{p=1}^{P} B_p(x) C_p(y).$$

From the linear equation $u = Aq$ and our construction of A, we then have

$$\begin{aligned} u_i &\approx \sum_{j=1}^{n} \sum_{p=1}^{P} B_p(x_i) C_p(y_j) q_j \\ &= \sum_{p=1}^{P} B_p(x_i) \left(\sum_{j=1}^{n} C_p(y_i) q_j \right) \\ &= \sum_{p=1}^{P} B_p(x_i) \hat{q}_p, \end{aligned}$$

where we defined $\hat{q}_p$ appropriately. Computing q is now reduced to two steps.

(1) Compute $\hat{q}_p = \sum_{j=1}^{n} C_p(y_j) q_j$ for $p \in \{1, 2, \cdots, P\}$. This costs $\mathcal{O}(nP)$.
(2) Compute $u_i = \sum_{p=1}^{P} B_p(x_i) \hat{q}_p$ for $i \in \{1, 2, \cdots, m\}$. The cost of this step is $\mathcal{O}(Pm)$. The total cost of these two steps is then $\mathcal{O}(P(m + n))$.

Aside: The matrix A has a low numerical rank, so its singular values decay exponentially. This is why we are able to take only P terms of the sum and achieve a precise approximation.

This process is easier to understand with an example. Returning to our problem from electrostatics, we let

$$G(x,y) = \begin{cases} \log(x-y) & x \neq y, \\ 0 & x = y. \end{cases}$$

If we choose a central point C_0 about which we calculate the potentials, we have $\log(x-y) = \log((x-c_0)-(y-c_0)) = \log(x-c_0) + \log(1-\frac{y-c_0}{x-c_0})$. For convenience, set $t := \frac{y-c_0}{x-c_0}$. Expanding this expression for the potential in a Taylor series, we have

$$\begin{aligned} \log(x-y) &= \log(x-c_0) + \log(1-\frac{y-c_0}{x-c_0}) \\ &= \log(x-c_0) + \sum_{p=1}^{\infty} \left(-\frac{t^p}{p}\right) \end{aligned}$$

For $|t| < 1$, the Taylor series is convergent. Because of this convergence (or, if you prefer, because the matrix A defined by G is of low numerical rank), we can truncate the series after obtaining P terms. This gives

$$\log(x-y) = \log(x-c_0) + \sum_{p=1}^{P} \left(-\frac{t^p}{p}\right) + E_p,$$

where the error term E_p satisfies

$$|E_p| \leq K \left(\frac{|x-C_0|}{|y-c_0|}\right)^{P+1}.$$